Online Appendix (Not for Publication)

A.1 Accounting Framework Derivation

This appendix provides the full derivation of the L-decomposition summarized in text (equation 3), including the commodity–industry distinction via the make and use tables.

A.1.1 Production and Materials

Each industry $i = 1, \dots, n$ has CES production as in the main text (equation 1). Following BEA's supply-use framework, industries may produce multiple commodities. Commodity output and material bundles are Cobb–Douglas aggregates:

$$Y_{ct} = \prod_{i=1}^{n} Y_{cit}^{\theta_{ci}}, \qquad m_{it} = \prod_{c=1}^{n'} \left[M_{cit}^{\zeta_{ci}} \tilde{M}_{cit}^{1-\zeta_{ci}} \right]^{\omega_{ci}}, \tag{19}$$

where M_{cit} and $\tilde{M}_{cit}$ denote domestically sourced and imported commodity c, θ_{ci} are commodity supply shares (make matrix Θ), ω_{ci} are input use shares (use matrix Ω), and ζ_{ci} are domestic sourcing shares, with $\sum_i \theta_{ci} = 1$ and $\sum_c \omega_{ci} = 1$. This is a first-order approximation—a CES aggregation would give identical results for log growth rates.

A.1.2 Cost Minimization and Growth Rates

Static cost minimization defines variable cost $C_i(Y_i) := \min_{N_i, m_i}\{W_i N_i + p_i m_i\}$ subject to (1), where p_i is the Cobb–Douglas price index of materials weighted by ω_{ci} and ζ_{ci}.

For any variable X_{it}, define $g_X^i := \log X_{it} - \log X_{it-1}$. Log-differencing the first-order and price-index conditions yields:

$$g_m^i = g_N^i - \rho\big(g_p^i - g_W^i\big), \tag{20}$$

$$g_p^i = \sum_c \omega_{ci}\zeta_{ci} \sum_{i'} \theta_{ci'} g_P^{i'} + \sum_c \omega_{ci}(1-\zeta_{ci})\, g_{\tilde{P}}^c, \tag{21}$$

$$g_{C/Y}^i \approx a_i g_W^i + (1-a_i) g_p^i - \Big(g_{Y/N}^i + \rho(1-a_i)(g_p^i - g_W^i)\Big). \tag{22}$$

Equation (20) is the log-differentiated CES first-order condition: materials demand tracks labor but falls when materials become relatively expensive, with the substitution response governed by ρ. Equation (21) decomposes material price growth into I–O-weighted domestic and import components. Equation (22) uses the log-linearized production function to substitute out TFP growth via $g_Z^i \approx \alpha_i(g_{Y/N}^i + \rho(1-a_i)(g_p^i - g_W^i))$ and the unit-cost identity $g_{C/Y}^i = g_C^i - g_Y^i$.[36]

The cost-share weight a_i is computed from observable labor shares:

$$a_i := \frac{\delta_i^{1/\rho}}{\delta_i^{1/\rho} + (1-\delta_i)^{1/\rho}} = \frac{ls_i}{ls_i + \sqrt{ls_i\,(1-ls_i)}}, \qquad ls_i := \frac{W_i N_i}{C_i}, \tag{23}$$

[36] Equations (20) and (21) are exact hat-algebra relations. Equation (22) is a first-order approximation around the symmetric point $N_i = m_i$ ($p_i = W_i$), though the approximation point is without loss of generality given the flexibility of choosing δ_i. Detailed derivations and linearizations are provided in the Mathematica notebook `decomposition.nb` in the replication package.

where ls is labor's share of variable costs, and $\delta_i = \left(1 + \left(a_i^{-1} - 1\right)^{\rho}\right)^{-1}$.

A.1.3 Margin Residual

Let $\mu_i := P_i Y_i / C_i(Y_i)$ denote the ratio of revenue to variable cost. Combining with (22) yields

$$g_{\mu}^{i} \approx g_{P}^{i} - a_i\, g_{W}^{i} - (1 - a_i)\, g_{p}^{i} + \underbrace{g_{Y/N}^{i} + \rho(1 - a_i)\big(g_{p}^{i} - g_{W}^{i}\big)}_{\text{variable factor productivity}}. \tag{24}$$

Under Cobb–Douglas production ($\rho = 1$), μ_i reduces to the standard markup of price over marginal cost (Hall, 1988; De Loecker et al., 2020); more generally it captures the joint effect of markups and returns to scale.

A.1.4 Stacking and the Leontief Inverse

Stacking (24) across industries and substituting the material price index (21) yields the linear system

$$g_{\mu} \approx L^{-1} g_P - B\, g_W + g_{Y/N} - D\big((\mathbf{1} - \zeta) \odot \Omega\big)^{\top} g_{\tilde{P}}, \tag{25}$$

where $L = (I - D\big(\zeta \odot \Omega\big)^{\top}\Theta)^{-1}$ is the Leontief inverse for cost-shock propagation and

- $\Omega = [\omega_{ci}] \in \mathbb{R}^{n' \times n}$: input *use* shares (from the BEA use table);
- $\Theta = [\theta_{ci}] \in \mathbb{R}^{n' \times n}$: commodity *supply* shares (from the BEA make table);
- $\zeta = [\zeta_{ci}] \in \mathbb{R}^{n' \times n}$: domestic shares of each commodity–industry pair;
- $D = \operatorname{diag}\big((1 - a_i)(1 - \rho)\big) \in \mathbb{R}^{n \times n}$: scales the materials channel by labor–materials complementarity;
- $B = \operatorname{diag}\big(a_i + (1 - a_i)\rho\big) \in \mathbb{R}^{n \times n}$: weights the wage contribution.

Premultiplying by expenditure weights w and inverting gives the aggregate L-decomposition:

$$g_P \approx \underbrace{\kappa^{\top} g_{\mu}}_{\text{Margins}} + \underbrace{\kappa^{\top} B\, g_W - \kappa^{\top} g_{Y/N}}_{\text{Unit labor costs}} + \underbrace{\kappa^{\top} D\big((\mathbf{1} - \zeta) \odot \Omega\big)^{\top} g_{\tilde{P}}}_{\text{Imports}}, \tag{26}$$

where $\kappa^{\top} := w^{\top} L$ are the Domar weights measuring each industry's total impact on consumer prices after all I–O linkages.

The key structural object is the Leontief inverse L, which propagates each industry's cost shocks through the full production network before aggregating with expenditure weights. In our implementation, w corresponds to PCE weights derived from the producer-price detail-level supply-use table. The elasticity ρ enters through both D (scaling network propagation) and B (scaling wage pass-through): lower ρ amplifies the contribution of material price shocks relative to wages.

A.1.5 Estimation of ρ

We estimate ρ using the CES first-order condition and the BEA/BLS KLEMS panel (63 industries, 1997–2023):

$$\log \frac{p_i m_i}{W_i N_i} = FE_i + FE_t + (1-\rho)\log \frac{W_i}{p_i} + \varepsilon_i. \tag{27}$$

The baseline estimate is $\hat{\rho} = 0.39$ (95% CI: $[0.06, 0.72]$), indicating significant labor–material complementarity. Differenced specifications yield lower point estimates (≈ 0.1–0.2), consistent with attenuation from measurement error in first differences.

A.2 Markup-Saturated Model

The baseline model presented in the main text sets $\mu_{i,t} = 0$ for all units and periods: the wall mechanism alone, together with exogenous prices and calibrated demand and productivity shifters, generates inflation structurally. This appendix presents a complementary specification—the *markup-saturated model*—in which exogenous markup wedges $\mu_{i,t}$ are added to each production unit's pricing equation and calibrated to close the remaining gap between model and data prices.

A.2.1 Setup

In the saturated model, the NKPC for each unit i is augmented with a time-varying markup shifter $\mu_{i,t}$ that enters the marginal cost expression as a multiplicative wedge:

$$\mathcal{MC}^{\mu}_{i,t} = \mathcal{MC}_{i,t} \cdot \exp(\mu_{i,t}). \tag{28}$$

The wedge $\mu_{i,t}$ is unit- and time-specific. It absorbs all forces that the structural model does not capture: market power, mismeasurement, demand-driven margin expansion, and any other residual between the model's predicted price and the data.

A.2.2 Calibration

The calibration proceeds in two stages. In the first stage, the baseline structural model is solved with $\mu_{i,t} = 0$: all time-varying parameters ($a_{i,t}$, $\xi_{i,t}$, τ_i) and the wall intensity γ_i are calibrated to match detrended employment, output, and productivity paths as described in Section 4. In the second stage, γ_i is held fixed at its baseline values and $\mu_{i,t}$ is introduced. The parameters $a_{i,t}$, $\xi_{i,t}$, and τ_i are then re-calibrated jointly with $\mu_{i,t}$ until the model matches all data targets—including GO deflator price paths for each tier. The iteration continues until convergence.

This two-stage procedure ensures that the wall mechanism is disciplined by micro evidence (stage 1) and not contaminated by the markup wedge (stage 2). The markup wedge is a residual that absorbs everything the wall does not explain. This approach is in the spirit of the business cycle accounting methodology of Chari et al. (2007), which introduces wedges into an otherwise standard model to diagnose which frictions are quantitatively important.

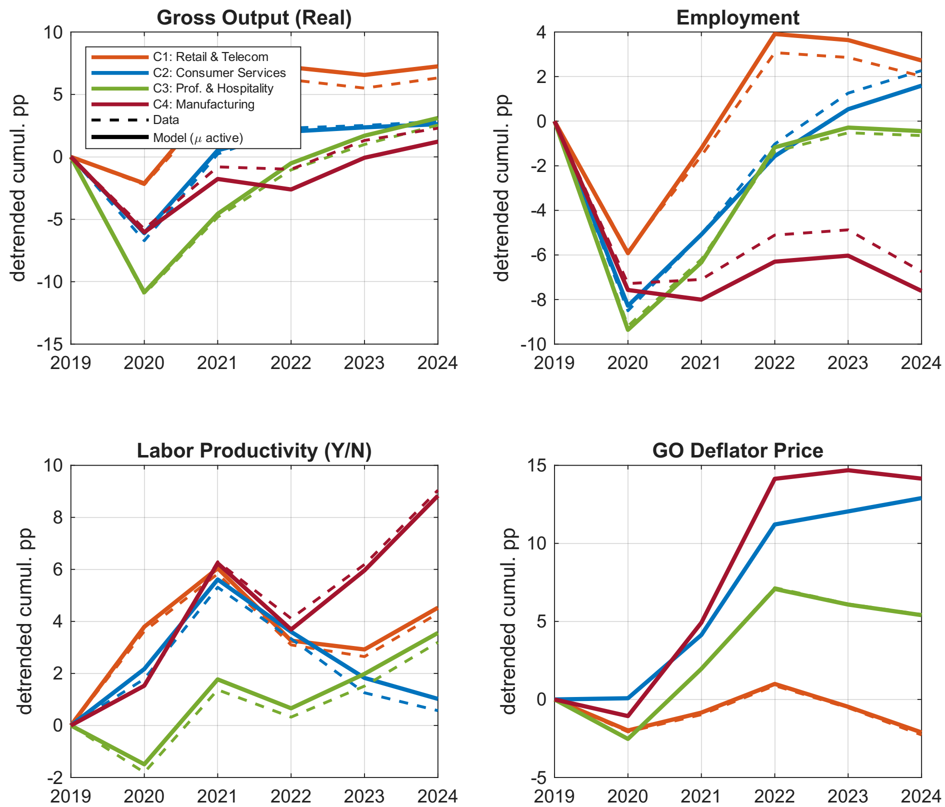

Figure 10: Chain-level fit: markup-saturated model. Gross output, employment, labor productivity, and GO deflators for each chain. Dashed: data. Solid: model with markup wedges active. All panels match by construction.

A.2.3 Fit

Figure 10 shows the chain-level fit of the saturated model. By construction, the model matches all data targets exactly—output, employment, productivity, and prices. The perfect fit confirms that the markup wedges successfully absorb all residual variation.

An advantage of the saturated specification is that agents form expectations over the *correct* equilibrium path—including the markup wedges that will prevail in each period. In the baseline structural model ($\mu = 0$), agents' expectations are formed over a path that underpredicts prices in some tiers, which may distort forward-looking pricing decisions. The saturated model eliminates this discrepancy, providing a cleaner counterfactual benchmark for decomposing the wall's contribution.

A.2.4 Counterfactual Decomposition

Table 11 reports the counterfactual decomposition. Each column reports the contribution of a specific channel, computed by setting that channel to zero and re-solving the full general equilibrium model. The wall contribution is Model $-$ CF($\gamma = 0$); the markup contribution is Model $-$ CF($\mu = 0$); and so on.

The wall remains the dominant channel, contributing +12.7 pp by 2022 and +15.3 pp by 2024. These values exceed the baseline structural estimate (+8.6 pp by 2022) because the saturated model's correct expectations amplify the

Table 11: Counterfactual Decomposition of Aggregate Inflation

Year	Model	Wall	Markup	Exo	Real Est.	GE
2020	-0.9	+2.1	+0.1	+0.1	+0.1	-3.2
2021	+2.3	+6.7	-0.4	+0.3	+0.2	-4.3
2022	+7.1	+12.0	+0.0	+0.8	+0.4	-5.8
2023	+7.4	+14.2	-1.2	+1.1	+0.6	-6.7
2024	+7.0	+13.7	-1.1	+1.5	+0.7	-7.1

Notes: Each counterfactual sets the named channel to zero and re-solves the full GE model. Wall = Model − CF($\gamma = 0$). Markup = Model − CF($\mu = 0$). Exo = Model − CF(HO/FO/EN= 0). Real Est. = Model − CF(HO= 0). GE = Model − (Wall + Markup + Exo). All values: N-PCE weighted cumulative log points, detrended.

wall's general equilibrium effects. Markup wedges contribute +1.9 pp by 2022 and +1.2 pp by 2024—a modest role relative to the wall. Exogenous prices (energy, real estate, finance) contribute +1.4 pp by 2022, growing to +3.0 pp by 2024 as energy and real estate effects accumulate. The large negative GE column (−7.8 pp by 2022) reflects the interaction between channels: shutting down each channel individually overstates its contribution because the channels interact nonlinearly in general equilibrium.

In the saturated model, the wall's share of total inflation (net of GE interactions) is approximately 87% by 2022 and 93% by 2024. These shares are larger than in the baseline because the markup wedges are modest and the wall's general equilibrium amplification is stronger under correct expectations.

A.2.5 Aggregate and Tier-Level Diagnostics

Figure 11 presents the aggregate price level under the saturated model alongside the structural ($\mu = 0$) counterfactual. The gap between the two lines measures the aggregate markup contribution. The structural model captures the direction and timing of the inflation surge but undershoots the level—the markup wedges close the remaining gap.

Figure 12 provides a comprehensive comparison of the saturated and structural specifications across six aggregate panels: prices, wages, unit labor costs, gross output, employment, and labor productivity. The real-side panels (output, employment, productivity) are virtually identical across the two models—the markup wedge enters the pricing equation but does not distort quantities. This confirms that the two-stage calibration procedure preserves the real allocation regardless of whether markups are active. The price panel shows the saturated model tracking data closely, while the structural model ($\mu = 0$) undershoots—the gap measures the aggregate markup contribution. The wage panel includes the ECI overlay, illustrating the gap between measured wages and effective hiring costs discussed in the main text.

A.2.6 Vertical Price Gradient and l-h Gap

Figure 13 shows the l-h price gap diagnostic under the saturated model. With markup wedges active, the model can match the data's vertical gradient across all four chains—including C_1's negative gap, which the structural model cannot reproduce (as discussed in the main text). The structural ($\mu = 0$) overlay confirms that the wall generates the correct sign for three of four chains but undershoots the magnitude.

The vertical-gradient validation in the main text (Section 5) uses the chain-level ℓ–h price gap rather than a tier-level V_k regression, for the reasons given in the footnote there. Figure 13 above carries the same comparison for

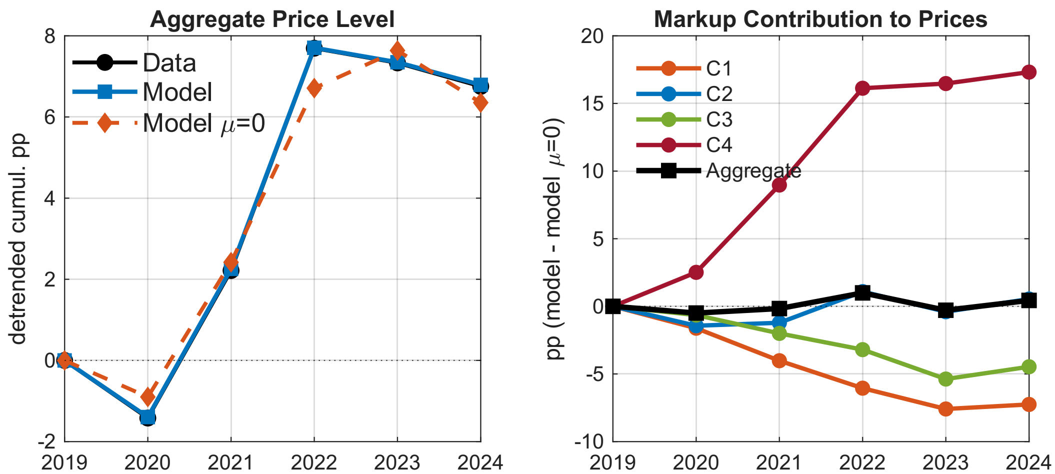

Figure 11: Aggregate inflation: saturated vs. structural model. PCE price level (cumulative log points vs. 2019, detrended). Data (black), saturated model with μ active (blue), structural model $\mu = 0$ (orange). The gap between the two model lines measures the aggregate markup contribution.

the saturated model: closing the per-tier price gaps with exogenous markup wedges pulls the model's ℓ–h gaps onto the data across all four chains, including C_1 where the $\mu = 0$ baseline alone cannot rationalize the sign.

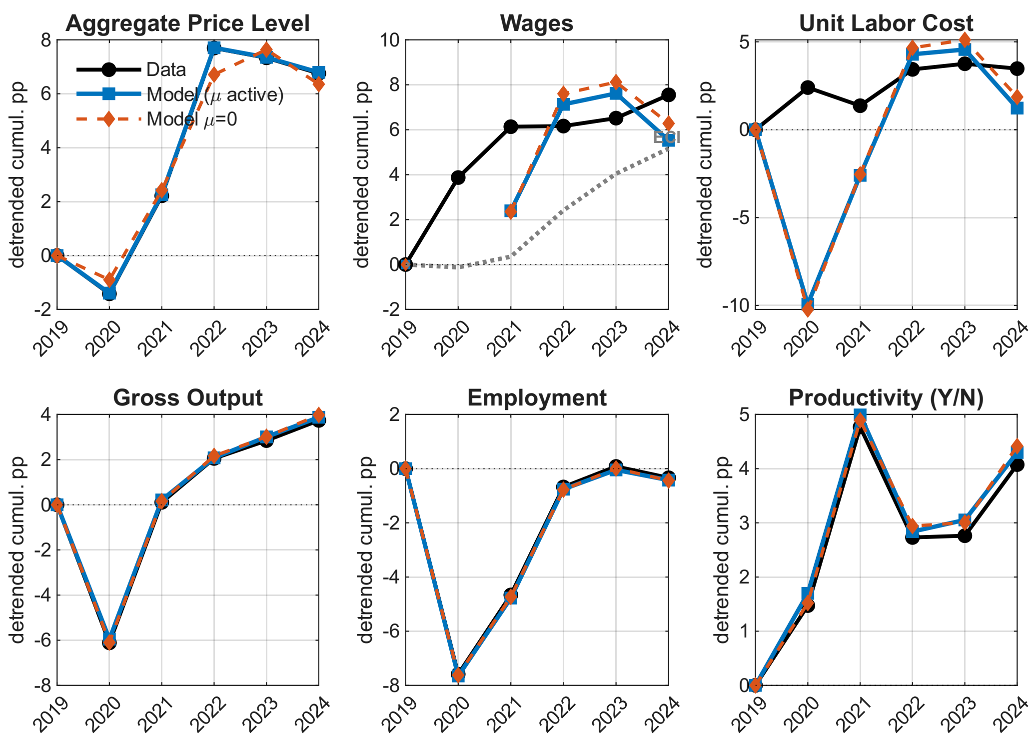

Figure 12: Aggregate fit: saturated vs. structural model. Six panels: prices, wages, unit labor costs (top); gross output, employment, labor productivity (bottom). Data (black), saturated model with μ active (blue), structural $\mu = 0$ (orange). Real-side panels are identical across specifications. ECI overlay on wage panel.

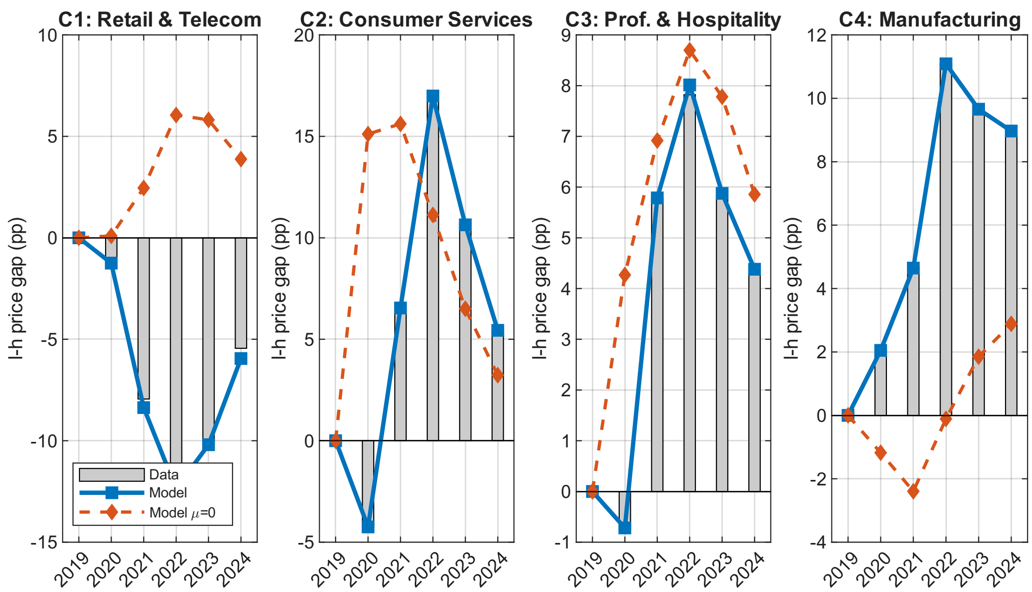

Figure 13: l-h price gap: saturated vs. structural. Cumulative upstream minus consumer-facing GO deflator gap by chain. Gray bars: data. Solid: saturated model (μ active). Dashed: structural ($\mu = 0$).

A.3 Robustness and Sensitivity

This appendix reports sensitivity analysis for the key model parameters. In each experiment, we fully recalibrate the time-varying parameters ($a_{i,t}$, $\xi_{i,t}$, τ_i) to match the same output and employment targets under the modified parameterization, holding γ_i fixed at baseline values. All results correspond to the baseline model configuration, where markups are held constant ($\mu = 0$) and the markup backfill mechanism is deactivated (as discussed here in the Online Appendix).

Upstream price stickiness. In the baseline, upstream (m/l) tiers have $\bar{\phi}_{ml} = 7$—near pass-through pricing that allows cost pressure to flow quickly through the supply chain. Raising upstream stickiness to $\bar{\phi}_{ml} = 12$ (equal to average h-tier stickiness) has a modest effect on aggregate inflation. Stickier upstream pricing creates a cost-pressure dam: wage increases accumulate in upstream marginal costs and arrive downstream as a larger, more persistent cost shock. A complementary experiment confirms that price stickiness and wall persistence (ψ) are complements, not substitutes: setting $\bar{\phi}_{ml} = 12$ but $\psi = 0$ collapses inflation to close to zero by 2022 and produces mild deflation by 2024. Without persistence, the wall fires once and dissipates—sticky prices slow the pass-through but there is nothing sustained to pass through.

Hiring cost premium. The baseline calibrates γ_i to a 35% peak effective hiring cost premium (within the 30–50% literature range). Scaling the peak target down to 30% (the lower bound, consistent with Davis and Haltiwanger 2022) reduces the wall's contribution roughly proportionally. The qualitative story is robust—the wall remains the dominant structural channel—but the quantitative magnitude depends on the assumed peak premium.

Downward price rigidity. Setting $\lambda = 0$ (symmetric Rotemberg costs) has a negligible effect on aggregate inflation. Removing downward rigidity produces deeper deflation in 2020 but the wall compensates during the recovery, leaving the net effect roughly neutral.

Elasticity of substitution. The labor–materials elasticity ρ_m matters modestly: setting $\rho_m = 0.95$ raises 2022 inflation by about +0.5 pp above baseline, while the wall contribution is essentially unchanged.

Expectations channel: backward indexation and discounting. Two experiments shut down the expectations channel through which the wall's cost impulse amplifies into aggregate inflation. The first sets $\iota_w = 0$ (full removal of backward wage indexation, down from the baseline $\iota_w = 0.55$): wages no longer respond to past inflation, breaking the wage–price feedback loop. Aggregate inflation falls from +5.00 pp to about +3.8 pp by 2022 and from +5.74 pp to about +3.6 pp by 2024; the wall contribution at 2022 drops from +8.37 to about +6.8 pp. The second experiment lowers the discount factor β well below the baseline ($\beta = 0.96$) so that firms attach much less weight to future marginal costs in their forward-looking Rotemberg pricing decisions. This delays and dampens the on-impact response: 2022 inflation drops to about +3.2 pp while 2024 inflation rises to about +8.2 pp, as accumulated cost pressure is released later rather than smoothed across periods. The two experiments tell a consistent story: the wall is necessary for inflation (without it the model deflates), but the aggregate magnitude and timing depend on how strongly wages and forward-looking expectations propagate the impulse. The wall is the impulse; wage–price indexation and forward-looking pricing are the amplifiers.

A.4 Counterfactual Experiment Implementation

This appendix describes the implementation of the counterfactual experiments reported in Section 6. All experiments use the baseline structural model ($\mu = 0$) and hold the wall intensity γ_i fixed at baseline values. The time-varying parameters ($a_{i,t}$, $\xi_{i,t}$, τ_i) are recalibrated under each modified specification to match the same (or modified) output and employment targets.

A.4.1 Experiment 1: `no_covid` — Removing the 2020 COVID employment collapse (the impulse)

This experiment shuts off the 2020 COVID shock by zeroing the 2020 detrended employment, output, and productivity targets for every tier, while leaving 2021–2024 cumulative targets unchanged at their data values. In the implementation (flag `CF_NO_COVID_2020 = true`) the modified targets are simply

$$\tilde{Y}_{i,2020} = 0, \qquad \tilde{N}_{i,2020} = 0, \qquad \widetilde{(Y/N)}_{i,2020} = 0,$$

for all 12 tiers, where the tildes denote the detrended (cumulative log) target relative to the 2019 steady state. The calibration then recovers a time-varying productivity/demand path ($a_{i,t}$, $\xi_{i,t}$, τ_i) that hits the 2021+ data levels without passing through the 2020 trough. The economy jumps from the 2019 steady state directly to the 2021 data level, and proceeds along the data path thereafter.

What is removed. The impulse: the concentrated 2020 employment loss and the subsequent reemployment flow associated with the 2020 lockdown year. In an aggregate plot of cumulative log N vs. time, the 2020 dip and the 2021 rebound are erased; the economy grows smoothly from 2019 to the 2021 data level instead.

What remains. The cross-tier asymmetry of the 2021+ recovery itself: chains that recovered quickly (e.g., C_2's middle tiers, which post strong positive detrended hiring in 2021–2022) still do so, and chains whose employment never returned to trend (e.g., C_4's upstream tiers) remain below trend. Walls therefore continue to fire asymmetrically during the 2021+ recovery—just without the additional impulse from the 2020 collapse-and-rebound cycle.

Recalibration. The initial wedge τ_i is unnecessary (no 2020 employment shock to match) and is set to zero. The productivity shifters $a_{i,t}$ and demand shifters $\xi_{i,t}$ are recalibrated via Newton iteration to match the modified targets. Because the 2020 collapse is absent, recovery hiring in 2021 is less concentrated: the wall fires ($\mathcal{W}_{i,t} > 0$) but at lower intensity, since the employment path rises smoothly from steady state rather than rebounding sharply from a deep trough.

Interpretation. The difference between baseline and `no_covid` inflation measures the contribution of the concentrated 2020 impulse itself, holding the 2021+ cross-tier recovery pattern fixed. By 2022, removing the 2020 COVID employment collapse lowers cumulative inflation by about +1.63 pp, or 33% of baseline inflation.

A.4.2 Experiment 2: `no_unevenness` — Removing cross-tier heterogeneity (supply + demand rotation)

This experiment shuts off cross-tier rotation. For each year t, every tier's detrended employment and output target is replaced with the Y_{ss}/N_{ss}-weighted aggregate across the 12 tiers. In the implementation (flag `CF_NO_ROTATION = true`), the recipe is applied directly to the detrended target matrices:

$$\overset{\downarrow}{Y}_{i,t} \leftarrow \sum_j w_j^{Y_{ss}} \overset{\downarrow}{Y}_{j,t}, \qquad \overset{\downarrow}{N}_{i,t} \leftarrow \sum_j w_j^{N_{ss}} \overset{\downarrow}{N}_{j,t}, \qquad \widetilde{(Y/N)}_{i,t} \leftarrow \overset{\downarrow}{Y}_{\mathrm{agg},t} - \overset{\downarrow}{N}_{\mathrm{agg},t},$$

for every tier i and year $t \in \{2020, \dots, 2024\}$. After the assignment all 12 tiers are identical and equal to the aggregate. The aggregate level—including the ≈ -7.6 pp 2020 collapse and the subsequent multi-year return toward steady state—is preserved exactly by construction.

What is removed. Cross-tier heterogeneity on both sides of the market: the supply side (uneven COVID employment destruction—some chains hit harder than others) and the demand side (the services-to-goods swing and its reversal). The two are zeroed simultaneously by this single recipe and cannot be separated by it; isolating supply rotation alone or demand rotation alone would require two additional experiments (`uniform_N_only` and `uniform_Y_only`), which we leave for future work.

What remains. The aggregate impulse: the 2020 collapse and its subsequent recovery are still in the model, just spread uniformly across all 12 tiers. Walls therefore fire everywhere, synchronously and at the aggregate rate; there is no within-chain or cross-chain cost gradient for the IO network to amplify.

Why aggregate-weighted, not equal-weighted. The procedure works in levels (not logs) and uses 2019 steady-state weights to avoid Jensen's inequality: setting all tiers to the same *log* growth rate would not preserve aggregate levels because $\sum_i Y_i \exp(g) \neq (\sum_i Y_i) \exp(g)$ when Y_i differ across tiers. The Y_{ss}/N_{ss}-weighted allocation ensures that aggregate output and employment in every year exactly match the baseline—only the cross-tier distribution changes.

Recalibration. With uniform targets, every tier experiences the same proportional shock and recovery. The time-varying parameters $a_{i,t}$, $\xi_{i,t}$, and τ_i are recalibrated via Newton iteration to match these modified targets. Walls fire uniformly across all tiers rather than concentrating in the tiers with the strongest recovery.

Interpretation. The difference between baseline and `no_unevenness` inflation measures the contribution of cross-tier heterogeneity (supply destruction plus demand rotation), holding the aggregate impulse fixed. By 2022, removing this cross-tier rotation lowers cumulative inflation by about +1.00 pp, or 20% of baseline inflation. The mechanism is that heterogeneity creates a vertical cost gradient—tiers that were hit hardest and recovered most generate disproportionate wall pressure—which the IO network amplifies into consumer prices.

A.4.3 Experiment 3: `no_covid_no_unevenness` — Removing both impulse and rotation

The combined experiment applies both flags simultaneously: `CF_NO_COVID_2020 = true` and `CF_NO_ROTATION = true`. There is no 2020 collapse for any tier, and the 2021+ paths are uniform across the 12 tiers. The residual aggregate path is smooth growth from the 2019 steady state to the 2024 data level, with every tier following the aggregate exactly. Walls still fire (positive aggregate hiring through 2021–2024) but minimally, because the path is smooth and symmetric.

Interpretation. The difference between baseline and `no_covid_no_unevenness` inflation measures the joint contribution of the two features of the COVID shock—the concentrated impulse and the uneven incidence—relative to a smooth, symmetric recovery. By 2022, the combined counterfactual produces only about +2.45 pp of inflation (vs. +5.00 pp at baseline); the two channels together therefore account for roughly +2.55 pp, or about 51% of baseline inflation at 2022. Broken out by Act, the joint contribution is sharply concentrated in Act I: the combined counterfactual accounts for about 89% of inflation in Act I (2019–2021) and only 29% in Act II (2021–2023), consistent with the two-act mechanism—Act I inflation is mostly driven by the COVID-specific pattern, whereas Act II is driven by the broader, economy-wide rebuild hiring firing capacity walls.

Why the shares are essentially additive in this calibration. Under the live γ-cap calibration, the two contributions add up nearly exactly: $1.00 + 1.63 = 2.63 \approx 2.55$, leaving an interaction term of only -0.08 pp (essentially zero). Earlier versions of the calibration exhibited a sizeable amplification effect between impulse and rotation (concentrated hiring in a narrow window pushing capped tiers further up the convex hiring-cost curve), but the current calibration's cross-tier cap on γ dampens that interaction. The economic point is unchanged: the two channels operate through the same wall mechanism, and a large interaction emerges whenever the wall's convexity bites at multiple tiers simultaneously—something the cap now prevents by construction.

A.4.4 Robustness Experiments

The robustness experiments summarized in Appendix A.3 modify structural parameters while holding γ_i fixed at its baseline value. Each experiment fully recalibrates $a_{i,t}$, $\xi_{i,t}$, and τ_i to the baseline output and employment targets under the modified parameterization. The model is then re-solved and the wall contribution is computed by the same homotopy procedure (setting $\gamma_i = 0$ and re-solving). Table 12 reports cumulative structural inflation in 2022 and 2024 and the 2022 wall contribution under each perturbation, against the baseline values of +5.00, +5.74, and +8.37 pp respectively.

We briefly describe each experiment and its main takeaway.

Upstream price stickiness ($\bar{\phi}_{ml} = 12$). The baseline sets upstream Rotemberg stickiness far below the h-tier level ($\bar{\phi}_{ml} = 7$), so cost pressure passes through the supply chain quickly. Raising $\bar{\phi}_{ml}$ to the h-tier average dams cost pressure in upstream tiers and releases it downstream with a lag. The effect on aggregate inflation is modest.

No persistence ($\psi = 0$, combined with $\bar{\phi}_{ml} = 12$). Setting wall persistence to zero kills the hysteresis channel: walls fire on current hiring but the premium decays fully within one period. Combined with sticky

Table 12: Cumulative inflation and wall contribution across robustness experiments.

Experiment	Π_{2022} (pp)	Π_{2024} (pp)	Wall_{2022} (pp)
Baseline	+5.00	+5.74	+8.37
$\bar{\phi}_{ml} = 12$ (stickier upstream)	+4.79	+5.61	+7.95
$\bar{\phi}_{ml} = 12$, $\psi = 0$ (no persistence)	+0.62	−0.54	+4.24
γ scaled to 30% peak	+1.07	+1.21	+4.76
$\lambda = 0$ (symmetric Rotemberg)	+5.39	+7.17	+10.77
$\rho_m = 0.95$ (labor–materials elasticity)	+7.18	+8.62	+9.11
$\rho_b = 0.95$ (materials-bundle elasticity)	+5.07	+5.78	+8.44
$\iota_w = 0$ (no wage indexation)	+3.79	+3.62	+6.83
β lowered (myopic pricing)	+3.22	+8.18	+5.47

Notes: Cumulative log points relative to 2019 steady state. Π_t is structural inflation aggregated at N-PCE weights; Wall_{2022} is the difference between the experiment's 2022 inflation and its $\gamma = 0$ homotopy counterpart. All experiments recalibrate the time-varying targets to match baseline Y and N paths under the modified parameterization.

upstream pricing, inflation collapses and the model turns mildly deflationary by 2024. This isolates the complementarity between persistence and stickiness: the wall's transient spike must be held long enough to propagate, and sticky pricing without persistence has nothing sustained to transmit.

Lower hiring-cost premium (γ_i scaled to 30% peak). Moving the peak effective premium from the baseline 35% to the literature low end (30%) shrinks wall contributions roughly proportionally—the wall contribution falls from +8.37 to +4.76 pp—and aggregate 2022 inflation drops to +1.07 pp. The qualitative finding—the wall is the dominant structural channel—is preserved, but quantitative magnitudes scale with the assumed peak premium.

Symmetric Rotemberg ($\lambda = 0$). Removing asymmetric downward price rigidity allows 2020 deflation to deepen and amplifies cumulative inflation in the recovery: 2022 inflation rises to +5.39 pp (vs. +5.00 baseline) and 2024 to +7.17 pp. The wall contribution rises to +10.77 pp as more pass-through becomes available in both directions.

Materials elasticity ($\rho_m = 0.95$). Pushing the labor–materials elasticity toward Cobb–Douglas matters: 2022 inflation rises to +7.18 pp (vs. +5.00 baseline) and 2024 to +8.62 pp. The labor–materials margin becomes more elastic, so wall-driven labor-cost shocks pass more strongly into prices. The qualitative conclusion is preserved, but the quantitative magnitude depends on this elasticity.

Materials-bundle elasticity ($\rho_b = 0.95$). Raising the elasticity *across* intermediate input varieties (away from the symmetric-CES baseline $\rho_b = 0.39$) has a near-zero effect: 2022 inflation is +5.07 pp (vs. +5.00 baseline). The materials-bundle CES is quantitatively irrelevant at the model's equilibrium price dispersion, so varying ρ_b in the [0.39, 0.95] range changes the wall contribution negligibly.

Zero wage indexation ($\iota_w = 0$). Shutting down backward wage–price feedback breaks the wage–price spiral: the wall still fires, but its cost impulse does not amplify through indexed wage growth. Aggregate inflation falls to +3.79 pp by 2022 and stabilizes near +3.62 pp by 2024. The wall is necessary for inflation (without it the model deflates), but the aggregate magnitude depends on how strongly wages propagate the impulse economy-wide.

Myopic pricing (β low). Lowering the discount factor substantially below the baseline 0.96 dampens the forward-looking channel in the Rotemberg pricing equation. Firms place less weight on future marginal costs, so the on-impact response is smaller (+3.22 pp in 2022) but accumulated pressure is released later (+8.18 pp by 2024) rather than smoothed across periods. Together with $\iota_w = 0$, this confirms the expectations channel as the amplifier: the wall is the impulse; backward indexation and forward-looking pricing set the magnitude and timing of the inflationary response.

A.5 Additional Tables

This section collects supplementary regression tables referenced in the main text.

Table 13 reports the robustness of the employment–inflation relationship (Fact 1) to using GO deflators instead of PPI as the price measure. GO deflators achieve full industry coverage by incorporating BEA imputations for service prices, but the results are attenuated relative to PPI, reflecting the inclusion of business service sectors where the physical supply chain mechanism is weaker.

Table 13: Reduced-Form Regressions: PPI vs. GO Deflators

	PPI		**GO (PPI sectors)**		**GO (all sectors)**	
	$\hat{\beta}$	p / t	$\hat{\beta}$	p / t	$\hat{\beta}$	p / t
Panel A: BEA-66 scatter — $\Delta\log N$ (19–20) $\to \Delta\log P$						
$\Delta\log P$ (19–23)	−1.21	$p = 0.038$	−0.10	$p = 0.773$	0.13	$p = 0.241$
$\Delta\log P$ (20–23)	−1.62	$p = 0.003$	−0.25	$p = 0.509$	0.09	$p = 0.413$
Panel B: Network depth — $Z_k \to$ cumulative inflation (19–22), BEA-66 FE						
Z_1		2.1**	−28	−1.0	−24	−0.9
Z_3	1370	7.3***	−549	−1.3	−289	−0.8
Z_5		—	−4428	−2.0**	−2036	−1.1
R^2 (Z_3)	0.81		0.38		0.57	
Panel C: Within-sector price dispersion — $\sigma^+(\Delta\log P)$ on Z_3						
$\hat{\beta}$	3.07	$t = 2.1$**	−7.37	$t = -1.4$	0.44	$t = 0.1$
BEA-66 sectors (Panel A)	20		26		58	

Notes: "PPI" uses producer prices from the BLS Producer Price Index (partial industry coverage). "GO (PPI sectors)" uses BEA gross-output deflators restricted to the same BEA-66 sectors that have PPI data. "GO (all sectors)" uses GO deflators for all 343 endogenous industries. Panel A: OLS at BEA-66 level. Panel B: GO-weighted WLS at BEA-357 with BEA-66 fixed effects. Panel C: OLS at BEA-66. GO deflators restricted to PPI-eligible sectors are directionally consistent with PPI but noisier: PPI detects the supply chain wall at Z_3 ($t = 7.3$), while GO deflators on the same sectors require deeper network averaging (Z_5: $t \approx 2$). Including all sectors eliminates the signal, confirming that the wall mechanism operates in goods-producing and physical-service sectors where input constraints bind, not in business services where IO linkages are organizational rather than physical. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 14 reports the full margin lifecycle regressions underlying the activation–deactivation pattern documented in Section 2.1.1. Network-exposed industries (Z_3) expand margins during Act I and compress them by 2022–2023, yet cumulative inflation remains strongly positive—consistent with the wall generating transient margin expansion that reverses as capacity is rebuilt.

Table 14: Capacity Wall Lifecycle: Z_3 and Margins vs. Inflation

	2019–2021 (Act I)	2021–2022	2019–2022 (Full)	2022–2023 (Act II)
Panel A: Margin residual				
$\hat{\beta}$	$+625.7^{***}$	$+556.9^{*}$	$+133.4$	-546.9^{**}
	(2.76)	(1.72)	(0.44)	(−2.17)
R^2	0.675	0.785	0.715	0.592
Panel B: Leontief-decomposed inflation				
$\hat{\beta}$	$+1074.6^{***}$	$+1692.8^{***}$	$+1408.9^{***}$	-185.1
	(6.87)	(7.86)	(7.32)	(−1.25)
R^2	0.659	0.855	0.810	0.660
n	356	356	356	356
BEA-66 FE	Yes	Yes	Yes	Yes

Notes: GO-weighted WLS regressions of margin residual (Panel A) and Leontief-decomposed inflation contribution (Panel B) on Z_3, the third-order IO network exposure to collapsed material suppliers. t-statistics in parentheses. Definitions of "collapsed" and "materials" as in Table 4. BEA-66 FE absorb between-summary-sector variation. All values in basis points of PCE contribution. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.6 Retail Profit Margins: Census QFS Evidence

This appendix documents the retail margin expansion discussed in Section 5 using data from the Census Bureau's Quarterly Financial Statistics (QFS, formerly QFR). The QFS provides quarterly income-statement data—revenues, operating costs, depreciation, and interest expense—for major industry groups, allowing direct measurement of operating margins at the sector level.

Operating vs. gross margins. We distinguish two margin concepts. The *operating margin* is the ratio of net sales to operating costs (excluding depreciation and interest), expressed as a percentage markup: $\mathrm{OM} = 100 \times (S/C_{\mathrm{op}} - 1)$. This is the narrowest measure, closest to the firm's pricing decision over variable costs. The *gross margin* additionally includes depreciation in the cost base: $\mathrm{GM} = 100 \times (S/(C_{\mathrm{op}} + D) - 1)$. Both measures are computed from annual aggregates (sum of four quarters) to smooth seasonal variation.

Retail trade. Figure 14 shows operating and gross margins for retail trade and its subcategories. All Retail Trade experienced a sharp margin expansion beginning in 2020, with the operating margin rising from approximately 7% to over 10%—a level not seen in the pre-pandemic sample. The expansion is broad-based: Food and Beverage Stores, General Merchandise and Clothing, and All Other Retail all show the same pattern, though with

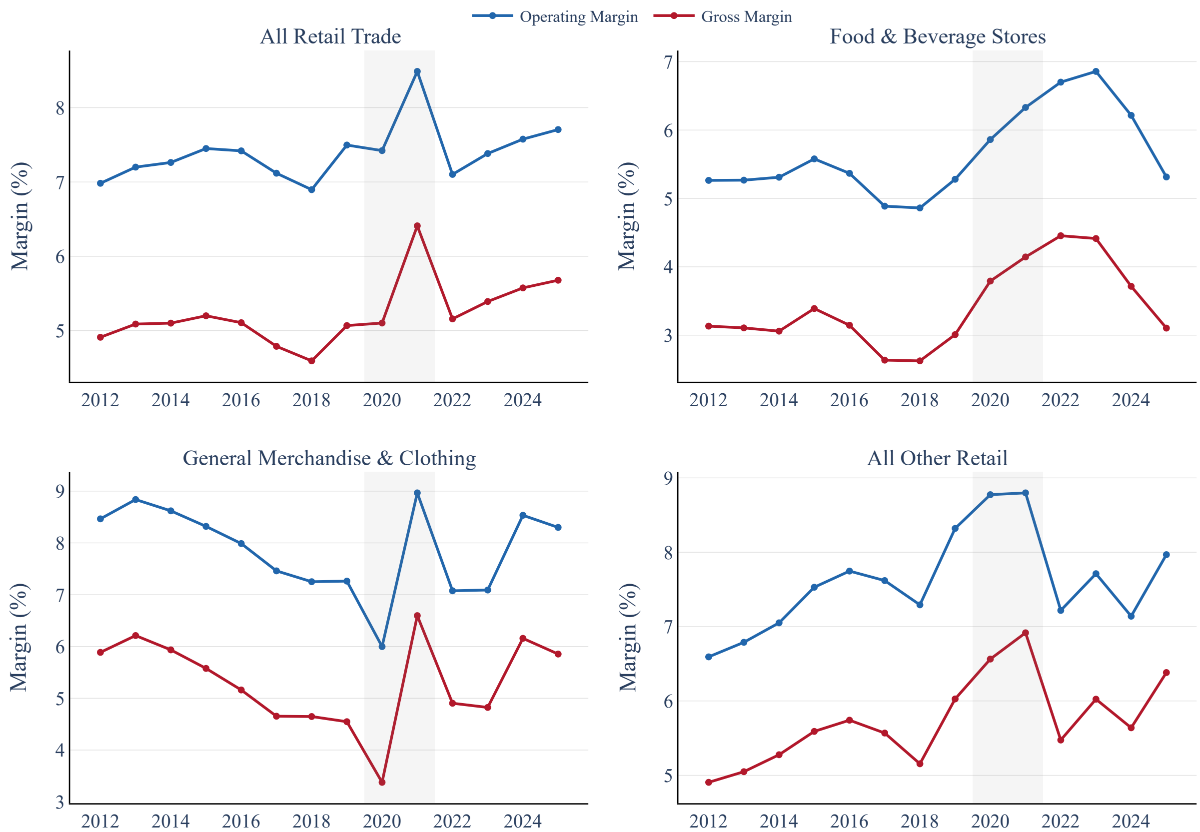

Figure 14: Retail trade operating margins by subcategory (Census QFS). Operating margin = $100\times(S/C_{\text{op}}-1)$; gross margin = $100\times(S/(C_{\text{op}}+D)-1)$. Annual (sum of quarters). Gray shading: COVID period (2020–2021).

different magnitudes. Food and Beverage margins rose the most in percentage terms, consistent with the well-documented divergence between food-at-home prices and upstream agricultural costs during the pandemic. The margin expansion persists through 2023, with only partial normalization by 2024.

Measurement challenges: retail margins across data sources. Comparing retail margins across datasets requires care because "retail output" is defined differently depending on the source. In the BEA input–output tables, retail output is the *trade margin*—the difference between the selling price and the purchase price of goods resold—so intermediate goods purchased for resale are netted out. Retail output in the IO tables is thus the value of the retail *service*, not total sales. The BEA/BLS KLEMS accounts follow the same national accounts convention: retail gross output equals the trade margin, and factor shares (labor, capital, intermediate inputs) are computed over this margin-based output concept. By contrast, firm-level data from Compustat or the Census QFS report standard income-statement items: total revenue includes all goods sold, and cost of goods sold (COGS) includes the purchase price of merchandise. The *gross margin* (revenue minus COGS) for a retailer is therefore close to the trade margin concept, but it excludes most labor costs—in retail, the vast majority of labor expenditure (store personnel, logistics, management) appears in selling, general, and administrative expenses (SG&A), not in COGS. The *operating margin* (revenue minus COGS minus SG&A) captures labor but also includes rent, marketing, and corporate overhead. No single income-statement line isolates the "cost of the retail service includ-

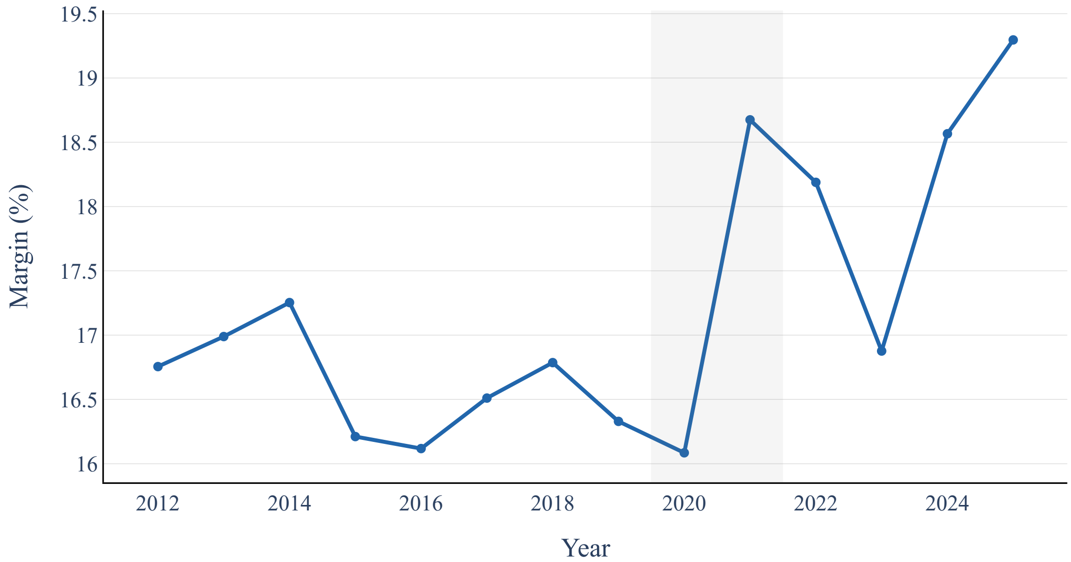

Figure 15: Operating margins: retail vs. rest of economy (Census QFS). "All Industry" excludes retail trade, petroleum and coal products, and primary metals. The retail margin expansion is sector-specific; non-retail margins remain within their historical range.

ing labor but excluding goods resold" that would be directly comparable to the BEA IO or KLEMS value-added concept. For this reason, we report QFS *operating* margins in this appendix: they provide the most complete cost base available from firm-level data, encompassing both merchandise costs and the SG&A expenses that contain retail labor. The margin expansion documented below is robust to the choice between operating and gross margin definitions, though the levels differ.

Non-retail industry. Figure 15 compares retail margins to the rest of the economy. The "All Industry" aggregate excludes retail trade, petroleum and coal products, and primary metals—sectors whose margins are dominated by commodity price swings rather than the supply chain dynamics we study. Outside retail, operating margins show no comparable expansion: they remain broadly flat through the pandemic and recovery, fluctuating within their historical range. This contrast is the key finding. The retail margin expansion is a sector-specific phenomenon, not a broad-based "greedflation" episode. Our structural model correctly identifies this: the capacity wall mechanism generates inflation in all four supply chains, but it cannot explain the retail margin expansion in C_{1h}, which requires an exogenous markup wedge (Section 5).

Retail-μ calibration target. The exogenous retail markup wedge in Section 5 (the C_1 line in Figure 9) is calibrated from the BEA Leontief decomposition applied to the C_1 consumer-facing tier. We use the PCE-weighted markup residual aggregated over the industries in C_1^h, excluding wireless telecommunications (WLTEL, the BEA wireless and other-telecom aggregate), whose 2022 establishment-level employment jump produces a spurious negative measured productivity and, as a mechanical consequence, a correspondingly large negative markup residual that is a classification artifact rather than a pricing decision. With WLTEL excluded, the retail-μ target is approximately +13.3 pp in Act I (2020–2021) and +6.7 pp in Act II (2022–2023). Because Rotemberg pricing frictions pass through only a fraction ($\approx$ 31%) of an injected markup wedge within the relevant horizon, the path

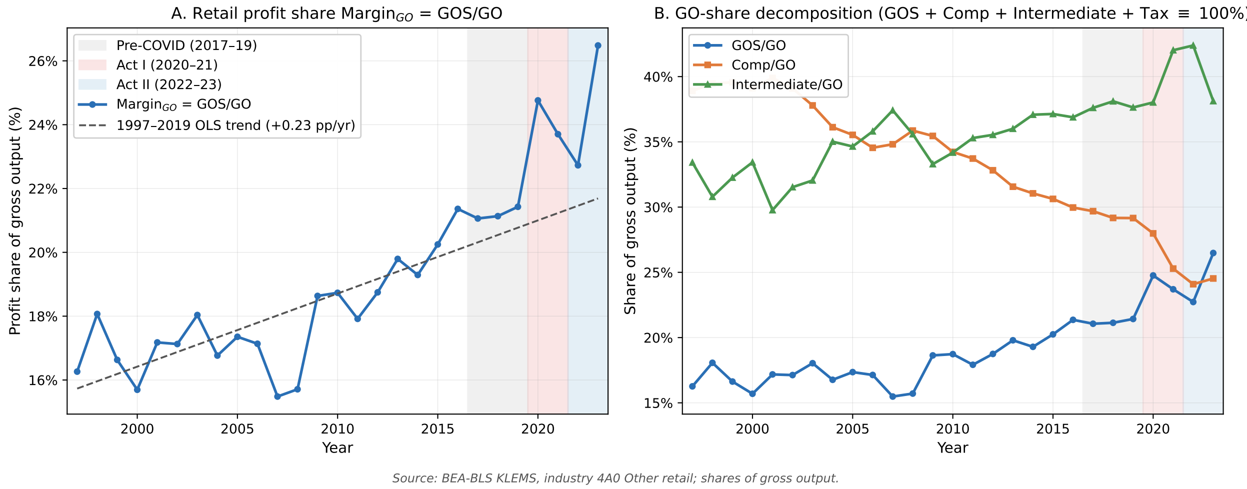

Figure 16: Retail profit share in BEA-BLS KLEMS (industry 4A0 "Other retail"). Left: annual gross-output profit share (GOS/GO) with 1997–2019 OLS trend ($\approx +0.23$ pp/yr); shaded windows mark Pre-COVID (2017–2019), Act I (2020–2021), and Act II (2022–2023). Right: decomposition of GO into GOS, compensation, and intermediate inputs as annual shares. Source: BEA-BLS KLEMS industry-level production account.

injected into the model is scaled by a factor of four relative to the raw data-derived target; this yields a realized model retail price that tracks the cumulative data path closely in Act I and overshoots moderately around 2022 due to the smoothing intrinsic to the Rotemberg mechanism.

KLEMS cross-validation (BEA-BLS industry-level production account). Figure 16 plots the KLEMS retail profit share for BEA Summary code 4A0 ("Other retail") from 1997 to 2023, with period averages summarized in Table 15. The dataset provides an independent level-based measure of retail margins that triangulates the markup wedge calibrated from our Leontief decomposition. We report the *gross-output profit share* (gross operating surplus divided by gross output, GOS/GO) rather than the more familiar value-added-based profit share (GOS/VA), so that the denominator matches our main-text framework's normalization by gross output. By the national-accounts identity GO = VA + intermediates, GOS/GO is mechanically smaller in level than GOS/VA but isolates the share of each dollar of gross output that becomes profit rather than compensation or purchased inputs. For the retail sector covered by KLEMS, the GO-based profit share averaged 21.2% over the pre-COVID window (2017–2019), rising to 24.2% in Act I (2020–2021; +3.0 pp) and 24.6% in Act II (2022–2023; +3.4 pp). The increase came primarily at the expense of the compensation share of gross output (which fell by roughly 5 pp over 2017–2023); the intermediate-input share rose temporarily during 2021–2022 and had largely reversed by 2023. Against a pre-COVID secular trend of about +0.23 pp per year (1997–2019), the post-COVID widening is roughly six to seven times the trend rate—a structural break rather than a continuation.

Comparing KLEMS to the model's markup wedge. KLEMS cannot be compared to the network-propagated markup contribution that the paper reports in Section 5, because the KLEMS accounts have no input–output structure. The natural comparison is instead with the paper's *single-pass* markup wedge—the own-industry residual before propagation through the Leontief inverse. Using the level-to-log-growth translation $d \ln M \approx d(\text{GOS}/\text{GO})/(1 - \text{GOS}/\text{GO})$, the KLEMS Act I shift of +3 pp in the profit share corresponds to roughly +4 log-% of markup-induced price wedge at the retail sector level. The single-pass μ we compute for the RTRADE aggregate over the same window is approximately +15 log-% cumulative. The two measures agree on direction

Table 15: Retail profit margin in BEA–BLS KLEMS (industry 4A0 "Other retail"), pre-COVID vs Acts I and II

Period	Margin$_{GO}$ (%)	Δ vs pre-COVID (pp)	Comp/GO (%)	Intermediate/GO (%)
Pre-COVID (2017–19)	21.2	—	29.3	37.8
Act I (2020–21)	24.2	+3.02	26.6	40.0
Act II (2022–23)	24.6	+3.40	24.3	40.3

Notes: Margin$_{GO}$ = Gross Operating Surplus / Gross Output, where GO = Value Added + Intermediate Inputs. Period averages over the three KLEMS years in each window. GO-based, not value-added-based, for alignment with the paper's Leontief framework. Share decomposition obeys the identity GOS/GO + Comp/GO + Tax/GO + Intermediate/GO ≡ 1; tax share not shown. Comparable to the paper's single-pass (own-industry) markup residual μ_i, not the L-propagated geo-pass, since KLEMS has no input–output structure. KLEMS 4A0 covers Detail codes 444, 446, 447, 448, 454 (building materials, health and personal care, gasoline stations, clothing, non-store retailers); BEA summaries 441 (motor vehicles), 445 (food and beverage), 452 (general merchandise) are not in this merged file. 2024 data unavailable (KLEMS release lag). Source: BEA–BLS KLEMS Industry-Level Production Account, 1997–2023.

and timing but differ in magnitude by roughly a factor of three, consistent with three considerations: (i) KLEMS 4A0 covers only a subset of retail (building materials, health and personal care, gasoline, clothing, and nonstore retailers; motor vehicles [441], food and beverage stores [445], and general merchandise [452] are not in the merged file), whereas our RTRADE is the broader 44–45 aggregate and captures precisely the subsectors most exposed to pandemic supply shocks; (ii) KLEMS 4A0 includes gasoline stations (447), which our pipeline classifies separately as part of the exogenous HO block; and (iii) our residual absorbs productivity-measurement noise and accrual-accounting timing differences that the KLEMS identities do not. KLEMS therefore corroborates the direction and two-act pattern of the retail markup widening while placing a conservative lower bound on its magnitude, consistent with the limits of any single data source.

Relation to the literature. Our retail finding sits between the two poles of the post-COVID markup debate. Glover et al. (2023) estimate that markups contributed a substantial share of 2021 inflation and then normalized by 2022–2023. Conlon et al. (2023), using firm-level data, find little-to-no link between markup changes and industry price increases, attributing the pandemic-era surge primarily to rising marginal costs. Our model is closer to Conlon for most of the economy—what appears in standard accounting as margin expansion is largely the wall-driven cost pass-through, not a change in pricing power—but closer to Glover specifically for retail, where an exogenous markup wedge is required in both acts to match the observed C_1 l-h price gap. Because our μ is a cumulating residual from annual price decompositions rather than a level deviation from marginal cost, magnitudes are not directly comparable to level markup estimates reported elsewhere in the literature. The "sellers' inflation" mechanism of Weber and Wasner (2023)—implicit coordination among firms in concentrated markets during periods of salient upstream cost pressure—remains a natural candidate microfoundation for the retail wedge we impose, consistent with the narrative framing in Section 1.

A.7 Data Sources

This appendix documents the primary data sources, derived inputs, and external series used across the empirical and structural components of the paper.

A. Primary Data Sources

1. **BEA Supply-Use Tables (2017 benchmark).** Detail-level supply-use tables from the Bureau of Economic Analysis, comprising 357 industries and approximately 389 commodities (harmonized to a 357×357 I-O system in our pipeline). We use the Use Table After Redefinitions (producer-price basis) to construct the IO matrix A_{jk}, commodity supply shares θ_{ci}, input use shares ω_{ci}, domestic shares ζ_{ci}, and PCE expenditure weights. Source:
`https://www.bea.gov/industry/input-output-accounts-data`.

2. **Quarterly Census of Employment and Wages (QCEW).** BLS establishment-level data aggregated to 6-digit NAICS, providing annual average employment, total annual wages, and annual average establishment counts. Coverage: 2017–2024, approximately 1,117 unique NAICS-6 codes. Used for: sector classification (establishment growth), within-sector employment and wage dispersion, shift-share instrument construction, and Phillips curve regressions. Source:
`https://www.bls.gov/cew/downloadable-data-files.htm`.

3. **Producer Price Index (PPI).** BLS output price indices at 6-digit NAICS, annual average observations. Coverage: 252 NAICS-6 industries with PPI series; coverage varies across the four supply chains, with the highest coverage in upstream manufacturing and consumer services tiers and the lowest in distribution (retail and wholesale) and professional-services tiers. Used for: within-sector price dispersion targets that identify ρ_s. Source:
`https://www.bls.gov/ppi/databases/`.

4. **BEA Gross Output Accounts.** Annual industry-level gross output, value added, and GO deflators from BEA GDP-by-Industry accounts. Used for: sector-level output, productivity, and Domar weight computation. Source:
`https://www.bea.gov/data/gdp/gdp-by-industry`.

5. **BEA/BLS KLEMS Dataset.** Joint BEA/BLS integrated industry-level accounts (63 industries, 1997–2023) providing labor compensation, intermediate input expenditure, and relative prices. Used for: estimating the labor–materials elasticity ρ via the KLEMS panel regression described in Section 2. Source:
`https://www.bls.gov/productivity/tables/`.

6. **BLS International Price Program.** Import price indices by end-use category, mapped to model sectors via BEA import commodity shares. Used for: exogenous import price shocks in the structural model. Source:
`https://www.bls.gov/mxp/`.

7. **NAICS Concordances.** Census Bureau crosswalks for 2017-to-2022 and 2022-to-2017 NAICS code mappings (approximately 1,150 concordance pairs). Used for: maintaining consistent industry panels across the 2022 NAICS revision. Source:
`https://www.census.gov/naics/`.

B. Model-Specific External Series

The structural model uses the following series directly:

1. **Employment Cost Index (ECI).** BLS ECI Total Compensation, private industry, by NAICS supersector. Provides a composition-free wage measure: unlike QCEW average wages (which shift with workforce composition during COVID), ECI tracks the same fixed basket of jobs over time. Used as the primary wage comparison for model validation. Annual log changes, detrended by 2017–2019 average annual growth. Source: `https://fred.stlouisfed.org/` (series CIU20XX).